

WHAT IS CLAIMED IS:

1. A network switch system for initializing a virtual volume in a system including a host system, and storage devices, the network switch system comprising:
 - a set of storage processors including first and second tier storage processors;
 - a Virtualization Coherency Manager (VCM) for receiving storage connectivity identifying which storage processors are connected to selected ones of the storage devices; and
 - a master storage processor for creating a logical tree based on the storage connectivity information, the logical tree reflecting a virtual volume of data distributed across the storage devices and includes (i) first tier objects representing partitions of the virtual volume data and (ii) second tier objects representing a logical configuration of the virtual volume,
 - wherein the VCM assigns the first tier objects to selective ones of the first tier storage processors and assigns the second tier objects to selective ones of the second tier storage processors based on the logical tree, and wherein each of the selective first and second tier storage processors initialize a local portion of the virtual volume to allow the host system to access the virtual volume through the network switch system.
2. The network switch system of claim 1, wherein the VCM designates a first tier storage processor as the master storage processor; and
 - the master storage processor collects the connectivity information from each storage device, wherein the connectivity information includes state information corresponding to virtual volume data included in the storage devices.

3. The network switch system of claim 1, wherein the master storage processor creates the logical tree by creating a system image of the virtual volume indicating which of the first tier storage processors are connected to particular ones of the storage devices storing virtual volume data associated with the virtual volume.

4. The network switch system of claim 1, wherein the logical configuration is a data protection configuration that distributes the virtual volume data across the storage devices.

5. The network switch system of claim 4, wherein the data protection configuration includes at least one of a striping configuration, a mirroring configuration, a concatenation configuration, and a striping over mirroring configuration.

6. The network switch system of claim 1, wherein the second tier objects include information referencing selective first tier objects based on the logical configuration of the virtual volume.

7. The network switch system of claim 1, wherein the master storage processor creates the logical tree based on virtual volume objects collected from selected ones of the storage devices, wherein the virtual volume objects identify which first tier storage processors are connected to corresponding ones of the storage devices.

8. The network switch system of claim 1, wherein the VCM is configured to assign a first tier object to a corresponding first tier storage processor based on whether the corresponding first tier storage processor is connected to a storage device including a partition of the virtual volume data.

9. The network switch system of claim 8, wherein the VCM is configured to assign the second tier objects by instructing the selected second tier storage processors to discover the first tier objects assigned to corresponding first tier storage processors.

10. The network switch system of claim 9, wherein the VCM is further configured to create a reference between a second tier object and a first tier object when the first tier object is associated with virtual volume partition data that is protected by the logical configuration provided by the second tier object.

11. The network switch system of claim 1, wherein each storage device includes a local portion of a virtualization state manager database (VSMDB) having virtualization objects associated with the local portion of the virtual volume, and wherein the VCM is configured to select a single first tier storage processor as the master storage processor and the master storage processor is configured to collect the virtualization objects stored in the VSMDB local portions included in the storage devices connected to the first tier storage processors.

12. The network switch system of claim 11, wherein the master storage processor is further configured to create the logical tree based on VSMDB data identifying which first tier storage processors are connected to particular storage devices having portions of the virtual volume data and send the logical tree to the VCM.

13. The network switch system of claim 1, wherein the master storage processor creates the logical tree by at least one of:

- logically partitioning the virtual volume data stored in the storage devices;
- logically mirroring the partitioned virtual volume data across multiple storage devices;
- logically striping the partitioned virtual volume data across multiple storage devices;
- logically striping the mirrored partitioned virtual volume data across multiple storage devices; and
- logically concatenating the partitioned virtual volume data across multiple storage devices.

14. The network switch system of claim 1, wherein each storage processor includes a Virtualization State Manager (VSM) and the VCM assigns the first tier objects by passing the first tier objects to the VSM of each first tier storage processor connected to a storage device storing a partition of the virtual volume.

15. The network switch system of claim 14, wherein the VCM passes selective ones of the second tier objects to the VSM of each second tier storage processor connected to the host system.
16. The network switch system of claim 15, wherein the VCM passes selective ones of the second tier objects to a first tier processor connected to the host system and to a storage device storing a partition of the virtual volume data.
17. The network switch system of claim 1, wherein the VCM assigns the second tier objects by assigning a second tier object to a corresponding second tier storage processor based on whether the corresponding second tier processor is connected to the host system.

18. A method for providing a virtual volume in a system including a host system, storage devices, and a network switch system having first and second tier storage processors, the method performed by the network switch comprising:
 - identifying, for each storage device, any first tier storage processors connected to the respective storage device;
 - creating a logical tree reflecting a virtual volume of data distributed across the storage devices, wherein the logical tree includes (i) first tier objects representing partitions of the virtual volume data and (ii) second tier objects representing a logical configuration of the virtual volume;
 - assigning the first tier objects to selective ones of the first tier storage processors;
 - assigning the second tier objects to selective ones of the second tier storage processors; and
 - managing the virtual volume using the first and second tier objects and storage processors.

19. The method of claim 18, wherein identifying includes:
 - designating a first tier storage processor as a Master Virtualization Storage Processor (MVSP); and
 - collecting, by the MVSP, virtualization objects from each storage device, wherein the virtualization objects include state information corresponding to virtual volume data included in the storage devices.

20. The method of claim 19, wherein creating the logical tree includes:
creating, by the MVSP, a system image of the virtual volume indicating which of the first tier storage processors are connected to particular ones of the storage devices storing the virtual volume data.
21. The method of claim 18, wherein the logical configuration is a data protection configuration that distributes the virtual volume data across the storage devices.
22. The method of claim 21, wherein the data protection configuration includes at least one of a striping configuration, a mirroring configuration, a concatenation configuration, and a striping over mirroring configuration.
23. The method of claim 18, wherein the second tier objects include information referencing selective first tier objects based on the logical configuration of the virtual volume.
24. The method of claim 18, wherein creating the logical tree includes:
creating the tree based on virtual volume objects collected from selected ones of the storage devices, wherein the virtual volume objects identify which first tier storage processors are connected to corresponding ones of the storage devices.
25. The method of claim 18, wherein assigning the first tier objects includes:

assigning a first tier object to a corresponding first tier storage processor based on whether the corresponding first tier storage processor is connected to a storage device including a partition of the virtual volume data.

26. The method of claim 25, wherein assigning the second tier objects includes:

discovering, by the second tier storage processors, the first tier objects assigned to corresponding first tier storage processors; and

creating a reference between a second tier object and a first tier object when the first tier object is associated with virtual volume partition data that is protected by the logical configuration provided by the second tier object.

27. The method of claim 18, wherein each storage device includes a portion of a virtualization state manager database (VSMDB) having virtualization objects associated with a local version of the virtual volume, and wherein identifying includes:

selecting a single first tier storage processor as a Master Virtualization Storage Processor (MVSP);

collecting, by the MVSP, virtualization objects stored in the VSMDB portions included in the storage devices connected to the first tier storage processors;

creating, by the MVSP, the logical tree based on VSMDB data identifying which first tier storage processors are connected to particular storage devices having portions of the virtual volume data; and

sending the system image to a Virtualization Coherency Manager (VCM).

28. The method of claim 27, wherein the VCM assigns the first and second tier objects to the first and second tier storage processors, respectively.
29. The method of claim 18, wherein creating the logical tree includes at least one of:
 - logically partitioning the virtual volume data stored in the storage devices;
 - logically mirroring the partitioned virtual volume data across multiple storage devices;
 - logically striping the partitioned virtual volume data across multiple storage devices;
 - logically striping the mirrored partitioned virtual volume data across multiple storage devices; and
 - logically concatenating the partitioned virtual volume data across multiple storage devices.
30. The method of claim 18, wherein the network switch system includes a Virtualization Coherency Manager (VCM) and each storage processor includes a Virtualization State Manager (VSM) and assigning the first tier objects includes:
 - passing, by the VCM, the first tier objects to the VSM of each first tier storage processor connected to a storage device storing a partition of the virtual volume data.
31. The method of claim 30, further including:

passing, by the VCM, selective ones of the second tier objects to the VSM of each second tier storage processor connected to the host system.

32. The method of claim 31, further including:

passing, by the VCM, selective ones of the second tier objects to a first tier processor connected to the host system and to a storage device storing a partition of the virtual volume data.

33. The method of claim 18, wherein assigning the second tier objects includes:

assigning a second tier object to a corresponding second tier storage processor based on whether the corresponding second tier processor is connected to the host system.

34. A method for accessing information in a system including at least one host system, storage devices, and first and second tier storage processors connected by a switching fabric, the method comprising:

identifying a first set of storage processors connected to the storage devices; collecting from each storage device connected to a first set storage processor, one or more objects associated with a virtual volume corresponding to a first host system;

creating a logical tree of the virtual volume based on the collected objects, wherein the logical tree includes first tier objects reflecting associations between the first

set storage processors and the virtual volume and second tier objects reflecting a logical configuration of the virtual volume;

assigning the first tier objects to selective ones of the first set storage processors; and

assigning the second tier objects to a second set of storage processors that are each connected to the first host system corresponding to the virtual volume,

wherein the first host system accesses data stored in the storage devices associated with the virtual volume through at least one of the second set storage processors.

35. A computer-readable medium including instructions for performing a method, when executed by at least one processor, for providing a virtual volume in a system including a host system, storage devices, and a network switch system having first and second tier storage processors, the method comprising:

identifying, for each storage device, any first tier storage processors connected to the respective storage device;

creating a logical tree reflecting a virtual volume of data distributed across the storage devices, wherein the logical tree includes (i) first tier objects representing partitions of the virtual volume data and (ii) second tier objects representing a logical configuration of the virtual volume;

assigning the first tier objects to selective ones of the first tier storage processors;

assigning the second tier objects to selective ones of the second tier storage processors; and

managing the virtual volume using the first and second tier objects and storage processors.

36. A computer-readable medium including instructions for performing a method, when executed by at least one processor, for accessing information in a system including at least one host system, storage devices, and first and second tier storage processors connected by a switching fabric, the method comprising:

identifying a first set of storage processors connected to the storage devices;

collecting from each storage device connected to a first set storage processor, one or more objects associated with a virtual volume corresponding to a first host system;

creating a logical tree of the virtual volume based on the collected objects, wherein the logical tree includes first tier objects reflecting associations between the first set storage processors and the virtual volume and second tier objects reflecting a logical configuration of the virtual volume;

assigning the first tier objects to selective ones of the first set storage processors; and

assigning the second tier objects to a second set of storage processors that are each connected to the first host system corresponding to the virtual volume, wherein the first host system may access data stored in the storage devices associated with the virtual volume through a second set storage processors.

37. A network switch system for providing a virtual volume in a storage virtualization environment including a host system, storage devices, wherein the network switch system includes first and second tier storage processors, the network switch comprising:

means for identifying, for each storage device, any first tier storage processors connected to the respective storage device;

means for creating a logical tree reflecting a virtual volume of data distributed across the storage devices, wherein the logical tree includes (i) first tier objects representing partitions of the virtual volume data and (ii) second tier objects representing a logical configuration of the virtual volume;

means for assigning the first tier objects to selective ones of the first tier storage processors;

means for assigning the second tier objects to selective ones of the second tier storage processors; and

means for managing the virtual volume using the first and second tier objects and storage processors.